The invention claimed is:

1. In a network of nodes connected to each other via bidirectional links, each of said nodes having a buffer for storing packets prior to transmission toward an ultimate destination, a method to control congestion on each of said links, said method comprising the steps of:

assigning a priority level λ_p from amongst at least two possible priority levels, to packets stored in a sending node X_ℓ buffer for transmission downstream via a link l to a receiving node R_ℓ, said link l being a portion of the path from said sending node X_ℓ to said ultimate destination;

transmitting upstream, via said link l, a feedback value f_ℓ from said receiving node R_ℓ to said sending node X_ℓ, said feedback value f_ℓ being indicative of the ability of said receiving node R_ℓ to store said packet in said receiving node R_ℓ buffer; and

transmitting downstream from said sending node X_ℓ to said receiving node R_ℓ, via said link l, only those packets stored in said sending node X_ℓ buffer whose priority level λ_p equals or exceeds the feedback value f_ℓ.

2. The method defined in claim **1** wherein said priority level λ_p is periodically changed when a packet is received in said receiving node R_ℓ, such that when a packet p with ultimate destination d arrives at R_ℓ from another network node (X_ℓ) over some link ℓ, the priority level λ^d of all packets at R_ℓ destined for node d, is updated as the maximum of

(a) the prior value of λ^d at R_ℓ, or

(b) $1+f_\ell$.

3. The method defined in claim **1** wherein the maximum value of said priority level λ_p is equal to the difference between (a) the maximum number D of nodes that a packet may

traverse through said network from any originating node to any ultimate destination, and (b) the number of nodes between said sending node X_ℓ and said ultimate destination node.

4. The method defined in claim **1** wherein said packets stored in said sending node X_ℓ buffer whose priority level λ_p equals or exceeds the feedback value f_ℓ are designated as eligible packets, and wherein said transmitting step includes processing said eligible packets in accordance with a prioritization algorithm.

5. The invention defined in claim **4** wherein said prioritization algorithm operates on a first-in/first out basis.

6. The invention defined in claim **4** wherein said prioritization algorithm operates on a round robin basis.

7. The invention defined in claim **1** wherein said feedback value f_ℓ is determined by

setting in the buffer at the receiving node R_ℓ thresholds B_i that limit the maximum amount of space for packets with priority levels λ^d less than or equal to i,

monitoring the priority levels λ^d of arriving and departing packets and the total space in the buffer at R_ℓ occupied by packets of various priority levels λ^d,

increasing priority levels λ_p of previously-stored packets, and

transmitting from the receiving node R_ℓ to the sending node X_ℓ a feedback value f_ℓ that represents the lowest priority level of packets that the receiving node R_ℓ could accept without violating any of the B_i buffer threshold constraints.

8. The invention defined in claim **7** wherein said increasing step includes periodically changing said priority level λ_p when a packet is received in said receiving node R_ℓ, such that when a packet p with ultimate destination d arrives at R_ℓ from another network node (X_ℓ) over some link ℓ, the priority level λ^d at R_ℓ associated with d is

updated as the maximum of

the prior value of λ^d at R_ℓ, or

$1+f_\ell$.

sub. a1

9. In a packet communication network comprised of interconnected nodes arranged to transmit variable length packets to adjacent nodes, wherein each node includes a buffer for storing packets enroute from a source node to a destination node, a method of controlling the transmission of a packet p from a sending node X_ℓ to a receiving node R_ℓ, via a link ℓ, said method comprising the steps of

sending from the receiving node R_ℓ to the sending node X_ℓ a feedback level f_ℓ such that there will be room in the buffer in the receiving node R_ℓ to store packets subsequently received from the upstream node X_ℓ;

assigning a priority level λ_p to packets stored in the buffer of the receiving node R_ℓ such that (a) all packets destined for the same destination have the same priority level, and (b) packets closer to their destination have a higher priority level; and

transmitting from the sending node X_ℓ to the receiving node R_ℓ, only those stored packets whose priority level λ_p is at least equal to the feedback level received from the receiving node R_ℓ.

10. The invention defined in claim **9** wherein D is the maximum number of hops that a packet must traverse through said network from a source one of said nodes to a destination one of said nodes, and wherein said assigning step includes assigning a level that is less than or equal to D minus the number of hops remaining between said receiving node R_ℓ and said destination.

11. In a packet communication network comprised of interconnected nodes arranged to transmit variable length packets to adjacent nodes, wherein each node includes a buffer for storing packets enroute from a source node to a destination node, a method of

controlling the transmission of a packet p from a sending node X_ℓ to a receiving node R_ℓ, via a link ℓ, such that (a) feedback is provided from each receiving node to each sending node regarding the fullness of the buffer at said receiving node, and (b) the occurrence of deadlocks and livelocks in said receiving node is avoided and no packets sent from said sending node X_ℓ to said receiving node R_ℓ are lost, said method comprising the steps of

transmitting from said receiving node R_ℓ to said sending node X_ℓ, a periodically updated transmit feedback parameter f_ℓ, said feedback value f_ℓ being determined by

(i) setting in the buffer at the receiving node R_ℓ thresholds B_i that limit the maximum amount of space for packets with priority levels λ^d less than or equal to i,

(ii) monitoring at the receiving node R_ℓ the priority levels λ^d of arriving and departing packets and the total space in the buffer at R_ℓ occupied by packets of various priority levels λ^d,

(iii) increasing priority levels λ_p of previously-stored packets, and

(iv) adjusting the feedback f_ℓ sent from the receiving node R_ℓ to the sending node X_ℓ to represent the lowest priority level of packets that the receiving node R_ℓ could accept without violating any of the B_i buffer threshold constraints,

assigning in said sending node X_ℓ, a level table associating, for each destination d to which said sending node may transmit a packet, a level λ^d, such that (a) λ^d is initially zero, (b) any packet in said node intended for destination d has the same level, and (c) when a packet arrives at sending node X_ℓ intended for destination d, λ^d is updated as the maximum of the previous value of λ^d or $(1+ f_\ell)$, whichever is greater, and

permitting sending node X_ℓ to send a packet to receiving node R_ℓ only if $\lambda^d \geq f_\ell$.

12. In a network of nodes connected to each other via bidirectional links, each of said nodes having a buffer for storing packets prior to transmission toward an ultimate destination, a method to provide feedback from receiving nodes to sending nodes to control packet transmission such that packets are not lost, and transmission of packets can occur

without creating overflow in said buffers and without creating deadlocks or livelocks, said method comprising the steps of:

assigning a priority level λ_p from amongst at least two possible priority levels, to packets stored in a sending node X_ℓ buffer for transmission downstream via a link l to a receiving node R_ℓ, said link l being a portion of the path from said sending node X_ℓ to said ultimate destination;

transmitting upstream, via said link l, a feedback value f_ℓ from said receiving node R_ℓ to said sending node X_ℓ, said feedback value f_ℓ being indicative of the ability of said receiving node R_ℓ to store said packet in said receiving node R_ℓ buffer;

transmitting downstream from said sending node X_ℓ to said receiving node R_ℓ, via said link l, only those packets stored in said sending node X_ℓ buffer whose priority level λ_p is at least equal to the feedback value f_ℓ; and

periodically adjusting said feedback value f_ℓ and said priority level λ_p.

Add a2